

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

THIS PAGE BLANK (USPTO)


PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶: H04Q 11/04	A1	(11) International Publication Number: WO 96/31081 (43) International Publication Date: 3 October 1996 (03.10.96)
(21) International Application Number: PCT/US96/03970 (22) International Filing Date: 22 March 1996 (22.03.96) (30) Priority Data: 08/410,712 27 March 1995 (27.03.95) US (71) Applicant: EMULEX CORPORATION [US/US]; 3535 Harbor Boulevard, Costa Mesa, CA 92626 (US). (72) Inventor: RICARD, Wayne, T.; 5225 Alamosa Park Drive, Oceanside, CA 92057 (US). (74) Agent: LAND, John; Fish & Richardson P.C., Suite 1400, 4225 Executive Square, La Jolla, CA 92037 (US).		(81) Designated States: AL, AM, AT, AU, AZ, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i> <i>With amended claims.</i>

(54) Title: WINDOW COMPARATOR**(57) Abstract**

A method and apparatus for reordering frames of data within an interface device of a receiving device. The interface device includes a window comparator (200) and a control device (202). The window comparator (200) includes a bidirectional 2:1 multiplexer (206), a counter (208), a first and second comparator (210, 212), a binary adder (214), a 3:8 decoder (216), eight 2-input OR gates (218), a subtractor (220), a bitmask (222), a shift register (224) with latched outputs, and a read buffer (204). The interface device can switch context to allow frames associated with more than one sequence to be handled concurrently. A "Missing Frame Window" is generated and associated with each sequence. When a frame associated with a particular sequence is received out of order, a bit within the Missing Frame Window representing the "missing" frame is asserted to represent receipt of the out of order frame. The out of order is then stored, and will later be transferred to a host.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LJ	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

WINDOW COMPARATOR

BACKGROUND OF THE INVENTION

1. *Field of the Invention*

This invention relates to a method and apparatus for digital communications, and more particularly to a method and apparatus for reordering sequential
5 information which has been received out of sequence.

2. *Description of Related Art*

In the computer industry, a number of protocols exist for moving mass quantities of data over high-speed digital networks. For example, Fibre
10 Channel is emerging as a standard for interconnecting workstations and peripherals within a computer system. Use of the Fibre Channel standard allows as much as 200 Mbytes of data to be transferred in as little as two seconds. Fibre Channel is considered by some to be a high-speed local-area network standard. Fibre Channel is a switched medium. This means that
15 packets of data are transferred by each of the users. Data is transported in packets of two kilobytes or smaller. These packet of data are referred to as "frames". Sequences comprise one or more frames which must be assembled in a predetermined order before the sequence can be considered complete at the receiving device.

20 FIGURE 1 is an illustration of the frame format for a Fibre Channel frame. The first four bytes of the frame are allocated to a "Start-of-Frame" field 101. The Start-of-Frame field 101 is a pattern of data that indicates that a frame follows. The next twenty-four bytes are allocated to a "Frame-Header" field 103. The Frame-Header field 103 includes four subfields that, taken together, uniquely
25 identify the frame. The first of these subfields is a source identification subfield (S-ID). The S-ID subfield is three bytes of information that identifies the device from which the frame originated. The second of these subfields is a destination identification subfield (D-ID). The D-ID subfield is three bytes of informa-

-2-

tion that identifies the device that is intended to receive the frame. The third of these subfields is a sequence identification subfield (SEQ-ID). The SEQ-ID subfield is three bytes that identify the particular sequence with which the frame is associated. The fourth of these subfields is a sequence count subfield (SEQ-CNT). The SEQ-CNT subfield determines the relative order of the frame within the sequence with respect to each of the other frames of that particular sequence.

The next 0 - 2112 bytes of the frame are allocated to the payload field (i.e., the field that carries the information that is to be transmitted from the source device to the destination device). The next four bytes of the frame are allocated to a cyclical redundancy code (CRC) field which permits errors in the data of the frame to be detected. The last four bytes of the frame are allocated to an "End-of-Frame" field.

One problem with communications protocols, such as Fibre Channel, is that data must be received in a predetermined order at the receiving end of a communications link. However, requirements of the communications link can prevent the data from arriving in order. For example, the source device may detect an error in the frame of information and transmit the next frame before correcting the error. After the error has been corrected, the frame that was in error is retransmitted out of order. Transmission of the next frame in order before correction of an error in a previously transmitted frame allows the source device to continue doing productive work (i.e., sending frames of data) while correcting a detected error. In another example, the source device may send frames in the proper order. However, frames may take different paths to the destination device. The result of different transmission delays over the different paths may result in the frames arriving out of order.

While transmission of subsequent frames after detecting an error in a previously transmitted frame, and transmission of frames over different paths, may cause frames to be received out of order, and thus burden the destination device with the task of reordering the frames, such operation allows the

-3-

source device to exploit the best delivery path and continue to transmit frames of data uninterrupted by errors.

Existing mechanisms for re-sequencing frames either penalize the source device by requiring manual or intelligent intervention at reduced performance, or penalize the system by requiring the sender to re-send the sequential information until it arrives at the destination in the correct order. For example, the TCP (Transport Control Protocol) Transport Layer protocol is a software mechanism that uses the services of Internet Protocol (IP) to provide reliable, ordered stream data delivery. The TCP software process in the receiving system checks each received segment to ensure that there are no missing, duplicated, or out of order segments. This system requires the receiver to intelligently intervene and thus reduces the performance of the receiver.

Accordingly, there is a need for a simple means for reordering frames received by a destination source to conform to the order intended by the source device which does not impact the efficiency and speed of the data communication.

SUMMARY OF THE INVENTION

The present invention is a method and apparatus for reordering frames of data within an interface device of a receiving device. Data is reordered to allow the receiving device to process the data in a predetermined sequential order. In accordance with the present invention, each frame that is transmitted from a source device is assigned a "Sequence Number" that indicates the relative position of that frame with respect to each other frame within a "Sequence" of data. A Sequence comprises one or more frames which must be assembled in a predetermined order before the Sequence can be considered complete at the receiving device. Each Sequence of data is also identified. The present invention can switch context to allow frames associated with more than one Sequence to be handled concurrently. A "Missing Frame Window" is generated and associated with each Sequence. Each Missing Frame Window

maintains a record of frames of a particular Sequence received out of order. When a frame of data associated with a particular Sequence is received, the Missing Frame Window associated with that Sequence is preferably loaded into a Missing Frame latch. The Missing Frame Window is preferably an 8-bit word. When a frame associated with a particular Sequence is received out of order, a bit within the Missing Frame Window representing the "missing" frame is asserted to represent receipt of the out of order frame. The out of order frame is then stored to be transferred to a host in order upon receipt of those frames which must be transferred before transfer of the out of order frame.

10 The least significant bit (LSB) of the Missing Frame Window represents the "next expected frame". The next expected frame is defined as a frame which carries a Sequence Number that has the lowest value within a range of possible values and which has not yet been received. When the LSB is asserted (indicating that the next expected frame has been received) the data
15 carried in the payload field of the next expected frame is transferred to the host. Each bit is shifted one bit toward the lowest order bit and the lowest order bit is again checked. This process repeats until the lowest order bit is not asserted.

20 The details of the preferred embodiment of the present invention are set forth in the accompanying drawings and the description below. Once the details of the invention are known, numerous additional innovations and changes will become obvious to one skilled in the art.

-5-

BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 is an illustration of the frame format for a Fibre Channel frame in accordance with the prior art.

FIGURE 2 is a block diagram of an interface device in accordance with the
5 present invention.

FIGURE 3 is a state diagram of the states of the state machine in accordance with the present invention.

Like reference numbers and designations in the various drawings refer to like elements.

10 DETAILED DESCRIPTION OF THE INVENTION

Throughout this description, the preferred embodiment and examples shown should be considered as exemplars, rather than as limitations on the present invention.

Overview

15 The present invention is an interface device that receives frames of data (i.e., multi-byte transmissions) over a communications link. Frames of data in accordance with the present invention preferably include the fields of information, such as shown in FIGURE 1 and described above. In the preferred embodiment of the present invention, each frame of data received
20 over the communications link in accordance with the present invention is associated with a "Sequence Number" carried in a field, such as the Frame-Header 103. For the sake of brevity, a frame associated with a Sequence Number "X" will be referred to as "frame X". For example, a frame associated with the Sequence Number "102" will be referred to as "frame 102".

-6-

5 The Sequence Numbers associated with each frame indicate the order in which the frames are intended to be received. For simplicity, frame "000" is considered to be the first frame of a Sequence in the present explanation, and each subsequent frame in order has a Sequence Number that is one greater than the previous frame in order. Therefore, frame "102" is intended to be the 103rd frame received. Frame "102" is intended to be received immediately after a frame "101". However, it should be clear that the range of possible Sequence Numbers may begin at any number and may increment at any predetermined interval.

10 The present invention transfers the information carried in the payload field 105 of each frame to a host in the order that each frame is intended to be received, as determined by the Sequence Numbers. If a frame is received out of order, then the data carried in the payload field 105 is stored. Only after transferring all the data of frames that should have been received prior, is the data of the out of order frame transferred. In one embodiment of the present invention, the host to which the data within the payload field 105 is transferred is a conventional computer. Alternatively, the host may be a computer peripheral device or any other device which serves as a destination node on a communications link.

20 In accordance with the present invention, the "next expected frame" is defined as a frame which carries a Sequence Number that has the lowest value within a range of possible values and which has not yet been received. For example, if the range of possible Sequence Numbers is 0 - 65,535 (as is the case in Fibre Channel frames), then if each frame from frame "000 to frame "100" has been received, and frame "104" has also been received, the next expected frame is frame "101", since "101" is the smallest Sequence Number not yet received, and yet within the possible range.

30 The present invention preferably stores eight bits (collectively referred to as a "Missing Frame Window") within a shift register with latched outputs. Each bit indicates whether any of the next frames within a predetermined range of

-7-

pected frames have been received. Alternatively, any memory means may be used to store the Missing Frame Window. The least significant bit (LSB) of the shift register is preferably associated with a next expected frame. Each other bit in order from the LSB to the most significant bit (MSB) is preferably associated with the next frame following in order. For example, if frame "003" is the next expected frame, then the LSB is associated with the frame "003". The second LSB is then associated with frame "004", the third LSB is associated with frame "005", etc. In an alternative embodiment, the MSB may be associated with the next expected frame, and thus the order of each bit is reversed.

Each bit changes state from a first logic state (deasserted) to a second logic state (asserted) when the frame associated with that bit has been received. When the LSB is asserted, the next expected frame has been received. Therefore, in accordance with one embodiment of the present invention, upon transition of the LSB, the most recently received frame is the next expected frame. Therefore, the most recently received frame can be transferred to the host, since the goal of the present invention is to transfer the frames to the host in the order expected. Once that frame has been transferred to the host, the shift register is updated with a value equal to the previous value shifted one bit toward the LSB. If the new value of the LSB is asserted, then the frame associated with that bit has been received and can be transferred to the host. If the new LSB is deasserted, then the interface device awaits receipt of the frame associated with the new LSB (i.e., waits for the LSB to be asserted).

In the preferred embodiment of the present invention, each of the parameters of the interface device may be stored in association with a particular Sequence and may then be reloaded upon receipt of a frame within that particular Sequence. This allows multiple information streams to be received and processed concurrently by the interface device.

Operational Description

-8-

FIGURE 2 is a block diagram of an interface device in accordance with the present invention. The interface device includes a window comparator 200 and a control device, such as a state machine 202. The state machine 202 controls the operation of the window comparator 200. The window comparator 200 includes a bidirectional 2:1 multiplexer 206, a counter 208, a first and second comparator 210, 212, a binary adder 214, a 3:8 decoder 216, eight 2-input OR gates 218, a subtractor 220, a bitmask 222, a shift register 224 with latched outputs, and a read buffer 204. The following signals are received in the state machine 202 as inputs which control the next state of the state machine 202: (1) the data received from incoming frames; (2) an MSB output signal; (3) a WRAP output signal; (4) a BORROW output signal; (5) an ERROR output signal; and (6) eight signals that represent the logical state of each bit of the shift register 224. FIGURE 3 is a state diagram of the states of the state machine 202.

The present invention starts upon receipt of a frame over the communications link. The state machine 202 enters INITIALIZATION state 301. When the frame is received, the state machine 202 processes the information contained in the Start-of-Frame field 101, the Frame-Header field 103, the CRC field 107, and the End-of-Frame field. The Start-of-Frame field preferably includes information that indicates whether this frame is the first frame of a Sequence. If the frame is the first frame of a Sequence, then the state machine 202 sets a first frame flag ("FF flag").

Upon detecting receipt of the Frame-Header 103, the state machine 202 preferably extracts a Sequence Identification (SEQ-ID) value and Sequence Counter (SEQ-CNT) value from the Frame-Header 103 and stores these values in a memory 230 within the state machine 202. In an alternative embodiment, the memory 230 may be external to the state machine 202. The SEQ-CNT is the Sequence Number of the frame. The SEQ-ID determines the particular Sequence with which the frame is associated. In an alternative embodiment of the present invention, the frame does not indicate the SEQ-ID.

-9-

In such an alternative embodiment, frames of a first Sequence should all be received before frame of a second Sequence can be received.

5 In addition, in INITIALIZATION state 301, each byte of information carried in the payload field 105 is separated from the overhead information carried in the other fields 101, 103, 107, 109 of the frame. The byte information is then preferably stored in a data field buffer within the memory 230. The length and the location of the buffer are preferably stored in the memory 230 in variables "DATA-LENGTH[x]" and "ADDR-PTR[x]", where "x" is a unique value associated with the particular frame (preferably by association with the SEQ-ID and SEQ-CNT associated with the frame).

10

The state machine then enters RESTORE state 302. In RESTORE state 302, the state machine loads the counter 208 and the shift register 224 with values associated with the SEQ-ID of the received frame. By allowing loading of the counter 208 and shift register 224 each time a new frame is received, the present invention ensures that the interface device can handle frames associated with a first Sequence interspersed with frames associated with at least a second Sequence.

15

If the FF flag is asserted, the state machine 202 loads both the shift register 224 and the counter 208 with the value of the Sequence Number associated with the first frame of the Sequence (preferably zero). If the FF flag is not asserted, then the state machine 202 loads the counter 208 and the shift register 224 with values that were previously saved in connection with the Sequence identified by the SEQ-ID. The values to be loaded into the counter 208 and shift register 224 are preferably identified by a table lookup indexed on the SEQ-ID. In an alternative embodiment of the present invention, the values to be loaded into the counter 208 and shift register 224 may be identified by use of a content-addressable memory, translation-lookaside buffer, or any other means for properly associating the values to be restored with the Sequence with which the last received frame is associated.

20

25

-10-

The state machine 202 controls a read/write signal coupled to the multiplexer 206 over a signal line 207 to cause the multiplexer 206 to receive input from the state machine 202 over signal line 209. A multiplexer select signal is coupled from the state machine 202 to the multiplexer 206 on a select signal line 211. The select signal determines whether the output from the multiplexer
5 is coupled to either the shift register 224 via the bitmask 222 or to the counter 208. The bitmask 222 is a two input multiplexer that determines whether the inputs to the shift register 224 are coupled to the multiplexer 206 or to the outputs of the OR gates 218. The value loaded in the shift register 224 is then
10 latched in the shift register 224. The value in the shift register 224 represents the Missing Frame Window and indicates whether a particular frame has been received. In an alternative embodiment, the OR gates 218 may have tri-state outputs.

In the preferred embodiment of the present invention, both the counter 208
15 and the SEQ-CNT are preferably 16 bits wide. The three least significant bits of the SEQ-CNT are coupled to a first input to the subtractor 220. The least significant three bits of the output from the counter 208 are coupled to a second input to the subtractor 220. The three least significant bits output from the counter 208 are subtracted from the three least significant bits from SEQ-
20 CNT of the first frame to aid in determining whether the most recently received frame is represented by a bit within the Missing Frame Window. It should be noted that more or less than the three least significant bits can be subtracted, depending upon the range of out of order frames that are to be acceptable. For example, if frames that are within a range of sixteen values relative to the
25 value of the next expected frame are acceptable, then the four least significant bits of the output from the counter 208 and from the SEQ-CNT would be coupled to the subtractor 220.

In the preferred embodiment, the three bits output from the subtractor 220 are coupled to the 3:8 decoder 216. The output of the decoder 216 is eight signal
30 lines, only one of which is asserted at any one time. The particular signal line to be asserted is determined by the state of the three signal lines output from

-11-

the subtractor 220 and coupled to the input of the decoder 216. Since there are eight possible combinations of logic states on the three input lines, each output line corresponds to a unique one of the eight combinations of the logic states possible on the three input lines. Accordingly, each output line from the decoder 216 represents a unique modulo-8 difference between the value of the SEQ-CNT signal and the value of the counter 208.

Each of the eight output signals from the decoder 216 is associated with one of the eight logical OR gates 218, one of eight inputs to the bitmask 222, one of eight outputs from the bitmask 222, one of eight inputs to the shift register 224, and one of eight parallel output lines from the shift register 224. Each of the eight output signals from the decoder 216 is coupled to the first input of the associated logical OR gate 218. The output of each OR gate 218 is coupled to the associated bitmask input. Each output from the bitmask 222 is coupled to the associated input to the shift register 224. Each output from the shift register 224 is then coupled to the second input to the associated OR gate 218. The particular decoder output associated with an input value of zero (i.e., the "zero" output) is coupled to the LSB of the bitmask 222, the "1" output is coupled to the bit of the bitmask 222 adjacent the LSB, and each other decoder output bit in order from the "2" to "7" is coupled to the bitmask input next in order. It should be noted that in RESTORE state 302, the bitmask input coupled to the multiplexer 206 is coupled to the input to the shift register 224.

Since the output of the decoder 216 represents the state of only the least significant three bits of the counter value and the SEQ-CNT value, the more significant output signals (i.e., bits 3-15) from the counter 208 and SEQ-CNT must also be taken into consideration in determining whether the most recently received frame is within the range of frames represented by the bits of the Missing Frame Window (i.e., within a predetermined range of values relative to the value of the next expected frame). The more significant output signals are coupled to the input of the binary adder 214 and to the first comparator 210. The first comparator 210 indicates that the higher order bits

-12-

output from the counter 208 and the higher order bits of the SEQ-CNT are equal by asserting the MSB signal at the output of the comparator 210.

In the preferred embodiment of the present invention, the output of the binary adder 214 is a 13-bit output that represents a binary value one greater than the value input to the adder 214. The output of the adder 214 is coupled to the input of the second comparator 212. The second comparator 212 compares the value represented by the more significant bits of the output of the counter plus one, with the more significant bits of the SEQ-CNT value. If these values are equal, then a WRAP signal output of the comparator 212 is asserted. The WRAP signal is used in combination with a BORROW signal output from the subtractor 220 to identify the case in which the value stored in the counter 208 is: (1) less than the value of the SEQ-CNT value (i.e., the Sequence Number of the most recently received frame is greater than the Sequence Number of the next expected frame) and (2) within a range of eight from the SEQ-CNT.

The combination of the MSB, BORROW, and WRAP signals provides a positive indication that the last frame received is associated with a Sequence Number that is greater than the Sequence Number of the next expected frame, and is within a predefined range relative to the value of the next expected frame (eight in the present example).

Therefore, if the MSB signal is asserted on the MSB signal line 217, then the most significant 13 bits of the counter 208 (which holds the value Sequence Number associated with the next expected frame) are equal to the most significant 13 bits of the SEQ-CNT (i.e., the Sequence Number associated with the last received frame). Accordingly, if the MSB signal is asserted, then the least significant three signals of SEQ-CNT must be equal to, or greater than, the three least significant signals of the Sequence Number associated with the next expected frame (the least significant bits of the value stored in the counter 208), as indicated by the BORROW signal being deasserted. If the BORROW signal is asserted, then the SEQ-CNT is in error (i.e., is less than the Sequence Number of the next expected frame, and thus out of the

-13-

predetermined range of values relative to the value of the next expected frame). Likewise, if the WRAP signal is asserted, then the three least significant bits of SEQ-CNT must be less than the three least significant bits of the value stored in the counter 208, as indicated by the BORROW signal being asserted. Otherwise, the most recently received frame is outside the range of valid frames and an error condition exits.

An additional error checking circuit 226 verifies that the last received frame is not associated with the same Sequence Number as a previously received frame. This is accomplished by comparing in eight discrete AND gates 228 each output from the decoder 216 with the associated output from the shift register 224. If the MSB signal is asserted, the BORROW signal is deasserted, and both an output from the decoder 216 and an associated output from the shift register 224 are asserted, then the last received frame is associated with the same Sequence Number as a previously received frame. Likewise, if the WRAP signal is asserted, the BORROW signal is asserted, and both an output from the decoder 216 and an associated output from the shift register 224 are asserted, then the last received frame is associated with the same Sequence Number as a previously received frame. In either of these cases, the state machine 202 enters an ERROR state 307. If the states of the WRAP, MSB, BORROW, and ERROR signals indicate that an error has occurred, the state machine 202 preferably asserts an interrupt request line to the host 201 to request assistance in dealing with the error condition. Alternatively, the state machine 202 may be intelligent enough to deal with error conditions without requesting assistance from the host 201. In either case, the particular action that is taken after detection of an error condition is outside the scope of the present invention.

If the state machine 202 determines that the SEQ-CNT value is within the predetermined range and is not equal to a previously received frame, the state machine enters LATCH state 303. In LATCH state 303, the state machine 202 causes the output signals from the OR gates 218 to be coupled to the parallel input of the shift register 224 through the bitmask 222 by controlling the select

-14-

signal to the bitmask 222. By ORing the previous value of the shift register 224 with the outputs of the decoder 216, those bits in the shift register 224 that were previously asserted remain asserted. Accordingly, if the value loaded into the shift register 224 in RESTORE state 302 had bits asserted to
5 indicate that frames had been previously received out of order, such information would be retained and supplemented by asserting the additional bit represented by the output from the decoder 216. The state machine 202 then causes the input signals to be loaded into the shift register 224. The output of the shift register 224 then indicates which frames within the
10 predetermined range of values relative to the value associated with the next expected frame have been received.

The state machine 202 then enters ADVANCE state 304. In ADVANCE state 304, the state machine 202 checks whether the LSB of the latch 204 is asserted. If the LSB of the latch 204 is asserted, then the last received frame
15 is the next expected frame. The state machine 202 then uses the variables DATA-LENGTH, and ADDR-PRT to identify the location of the data buffer for the byte information stored in memory and associated with the counter value (i.e., the last received frame). The byte information is recovered and then transferred to the host 201.

20 The state machine 202 increments the counter 208 by one and causes the contents of the shift register 224 to be shifted one bit toward the LSB. If the new LSB is asserted, then the state machine remains in ADVANCE state 304 and uses the variables DATA-LENGTH, and ADDR-PTR to identify the location of the byte information stored in memory and associated with the new
25 counter value. The byte information is recovered and then transferred to the host 201. The counter 208 is again incremented and the value of the shift register 224 is again shifted one bit toward the LSB. This process continues until the LSB is zero. If the new value of the LSB is zero, then the state machine 202 enters SAVE_CONTEXT state 305. In SAVE_CONTEXT state
30 305, the values of the Counter 208 and shift register 224 are saved in the memory 230. The state machine 202 preferably receives the value from the

-15-

shift register 224 through the buffer 204 and multiplexer 206. In addition, the current contents of DATA-LENGTH, and ADDR-PTR are also saved.

The state machine 202 then returns to INITIALIZATION state 301 and awaits receipt of the next frame. When a next frame is received, the SEQ-ID and SEQ-CNT are extracted to determine the address of the context to be restored to the counter 208 and shift register 224 (i.e., the values that were last maintained in the counter 208 and shift register 224 in connection with the last frame that was associated with the same Sequence). The byte information carried in the payload field 105 is buffered in memory at a location identified by the variables DATA-LENGTH, and ADDR-PTR (each of which is associated with the SEQ-CNT of this frame) until it is determined that this frame is the next expected frame (and consequently, the counter value is equal to the SEQ-CNT of this frame). The new contents of the shift register 224 is then loaded into the latch 204. The comparators 210, 212 of the Window Comparator 200 then provide the logic state of the WRAP, BORROW, and MSB signals based upon the value of the SEQ-CNT and the new counter value.

If these signals indicate that the SEQ-CNT is within the range of values relative to the value of the next expected frame, then the state machine 202 enters ADVANCE state 304 once again. The LSB of the latch 204 is once again checked. If the new LSB is asserted, then the state machine 202 uses the variables DATA-LENGTH, and ADDR-PTR associated with the value of the counter 208 to identify the location of the byte information stored in memory. The byte information is recovered and then transferred to the host 201. The state machine 202 again increments the counter 208 and shifts the contents of the shift register 224 one bit toward the LSB. If the new LSB is asserted, the variables DATA-LENGTH, and ADDR-PTR associated with the value of the counter 208 identify the location of the byte information stored in memory. The byte information is recovered and then transferred to the host 201. The state machine 202 again enters ADVANCE state 304, increments the counter 208, and shifts the contents of the shift register 224 one bit toward the LSB. This process is repeated until the new LSB is not asserted.

-16-

When LSB of the latch 204 is not asserted, the last received frame was not the next expected frame. The state machine 202 then enters SAVE_CONTEXT state 305 and saves the values of the shift register 224 and the counter 208, as well as the values of the DATA-LENGTH, and ADDR-PTR. Next, the state machine 202 enters INITIALIZATION state 301 and awaits the receipt of the next frame. This procedure is continuously repeated.

The present invention has several advantages over the prior art. The implementation of the present invention allows the invention to be embodied using simple logic elements, making it suitable for embedding in digital electronic communications systems. Furthermore, the present invention provides a simple mechanism to allow a higher level control function to determine the state of the interface device and force an alternate state. This facilitates "context switching", or the parallel use of a relatively complex machine by alternatively saving and restoring simple state information (i.e., the values of the counter and shift register). In addition, if frames have a field that identifies the particular Sequence, the present invention is capable of multiplexing between multiple streams of data by updating the values of the counter and shift register with values relevant to the Sequence of the last received frame. Because the present invention requires very little processor intervention, there is very low latency between the receipt of the next expected frame and the transfer to the host of the data carried in the payload field of that frame. If frames are received by the interface device in sequential order, then there is very little delay in the transfer of the payload information to the host for each frame. Since the Missing Frame Window identifies a predefined range of frames that may be received, an indication that a frame is out of the range allows the invention to indicate when temporary storage is, or would be, overburdened. The present invention also indicates an error condition when the same frame has been received twice. Still further, the present invention easily accommodates transmission protocols in which the length of the received information varies between deliveries and where the nature of the variations is unpredictable by the receiver.

-17-

A number of embodiments of the present invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, the present invention is described in the context of a state machine which controls
5 the operation of the window comparator. However, any circuitry, such as discrete logic, programmed general purpose computer, or programmable logic array may be used to perform the functions of the state machine described above. Likewise, the window comparator circuit of the present invention can be implemented in any means, such as, programmed general purpose
10 computer, programmable logic array, or state machine. Furthermore, the number of bits (i.e., the range of the Missing Frame Window) may vary from the 8-bits described in the exemplar. Accordingly, it is to be understood that the invention is not to be limited by the specific illustrated embodiment, but only by the scope of the appended claims.

CLAIMS

1. An interface device for coupling a receiving device to a communications link and for ensuring that frames of data received by the interface device, each frame being included within a sequence of frames, are communicated to the receiving device in a predetermined order defined by an external transmitting device, the interface device including:
 - (a) a control device; and
 - (b) a comparator window circuit coupled to, and controlled by, the control device, the comparator window circuit having a missing frame window storing means for storing a missing frame window which indicates whether the last received frame is the next expected frame in the order defined by the external transmitting device.
2. The interface device of claim 1, wherein received frames include information that identifies a sequence in which the frame is included, and the comparator window further includes a loading means, coupled to the missing frame storing means, for loading the missing frame window storing means with a missing frame window associated with the sequence of the last received frame.
3. The interface device of claim 1, wherein the control device transfers data received by the interface device within a select frame only upon determining from the missing frame window that the select frame is the next expected frame.

4. The interface device of claim 1, wherein each valid frame received by the interface device is associated with a value that indicates the relative position of that frame within the order, and the comparator window circuit further includes:
 - (a) a storage means for storing the value associated with the next expected frame; and
 - (b) a comparison means, coupled to the storage means and the control device, for comparing a value associated with the next expected frame with a value associated with a last received frame and indicating whether the value associated with the last received frame is within a predetermined range of values relative to the value of the next expected frame.
5. The interface device of claim 4, wherein the control device transfers data received by the interface device within a select frame only upon determining from the comparison means and the missing window frame that the select frame is currently the next expected frame and the select frame was associated with a value within the predetermined range of values at the time the select frame was received.
6. The interface device of claim 5, wherein the storage means is a counter operatively coupled to the control device and capable of incrementing the value stored therein by a predetermined value upon command from the control device.
7. The interface device of claim 6, wherein the counter is commanded by the control device to increment by the predetermined value when the missing frame window indicates that the last received frame is the next expected frame.

8. The interface device of claim 7, wherein the window comparator circuit further includes a shift register coupled to the control device and coupled to an input to the missing frame window storing means, such that a missing frame window input to the missing frame window storing means may be shifted under the control of the control device prior to being stored in the missing frame window storing means when the missing frame window indicates that the last received frame is the next expected frame.
9. The interface device of claim 7, wherein: (1) the control device has value signal lines that output the binary value associated with the last received frame, (2) the counter has output signal lines that output the binary value of the counter, and (3) the comparator window circuit further includes:
 - (a) a first comparator having an output and a first and second input, the first input coupled to the output signal lines of the counter, the second input coupled to the value signal lines, and the output of the comparator coupled to the control device, the first comparator for determining when the value associated with the last received frame is within a predetermined range of values relative to the value of the counter;
 - (b) an adder having an input and an output, the input coupled to the output signal lines of the counter, for outputting a value that is greater than the value of the counter; and
 - (c) a second comparator, having an output and a first and a second input, the first input coupled to the output of the adder, the second input coupled to the output signal lines of the counter, and the output of the second comparator being coupled to the control device, the second comparator for comparing the output of the adder to the value of the counter and providing the result of the comparison to the control device;

-21-

- (d) a subtractor having an output, a first input, and a second input, the first input coupled to at least one of the least significant output lines from the counter, and the second input coupled at least one of the least significant output lines from the control device, for representing in binary form, the difference between the value represented on the least significant lines of the counter and the value represented on the least significant lines from the control device;
- (e) a decoder having an input and an output, the output including a plurality of output signal lines, each output signal line being associated with a unique binary input value coupled to the decoder from the subtractor; and
- (f) a plurality of two-input OR gates, each OR gate being associated with a unique one of the decoder output signal lines, a first input of each OR gate being coupled to the associated decoder output signal line, the output of each OR gate being coupled to a unique input to the missing frame window storing means for storing values indicative of the values associated with frames received by the interface device.

10. The interface device of claim 7, wherein the comparator window circuit further includes:
- (a) means for determining whether a predetermined number of high order bits of the value associated with the last received frame are equal to or one greater than an equal number of high order bits of the value associated with the next expected frame;
 - (b) means for asserting one of a plurality of difference signal lines, each difference signal line associated with a unique value, when the difference between the least significant bits of the value associated with the last received frame and an equal number of the least significant bits of the value associated with the next expected frame is equal to the unique value associated with the difference signal line to be asserted; and
 - (c) a plurality of two-input OR gates, each OR gate being associated with a unique one of the plurality of difference signal lines, a first input of each OR gate being coupled to the associated difference signal line, the output of each OR gate being coupled to a unique input to the missing frame window storing means for storing values indicative of the values associated with frames received by the interface device, each value being within a predetermined range of values relative to the next expected frame, the second input of each OR gate being coupled to an output from the missing frame window storing means.

AMENDED CLAIMS

[received by the International Bureau on 26 August 1996 (26.08.96);
original claims 1-10 replaced by amended claims 1-6 (5 pages)]

1. An interface device for coupling a receiving device to a communications link and for ensuring that frames of data received by the interface device, each frame being included within a sequence of frames, are communicated to the receiving device in a predetermined order defined by an external transmitting device, the interface device including:
 - (a) a control circuit for processing information contained in each frame of data; and
 - (b) a comparator window circuit coupled to, and controlled by, the control circuit, wherein each valid frame received by the interface device is associated with a value that indicates the relative position of that frame within the predetermined order, the comparator window circuit having:
 - (1) a missing frame window storing means for storing a missing frame window which indicates whether the last received frame is the next expected frame in the order defined by the external transmitting device;
 - (2) a counter, operatively coupled to the control circuit, for storing the value associated with the next expected frame and capable of incrementing the value stored in the counter by a predetermined value upon command from the control device when the missing frame window indicates that the last received frame is the next expected frame,

- 24 -

- (3) a comparator, coupled to the counter and the control circuit, for comparing a value associated with the next expected frame with a value associated with a last received frame and indicating whether the value associated with the last

received frame is within a predetermined range of values relative to the value of the next expected frame, and

- (4) a shift register coupled to the control circuit and coupled to an input to the missing frame window storing means, such that a missing frame window input to the missing frame window storing means may be shifted under the control of the control circuit prior to being stored in the missing frame window storing means when the missing frame window indicates that the last received frame is the next expected frame.

- 2. The interface device of claim 1, wherein received frames include information that identifies a sequence in which the frame is included, and the comparator window circuit further includes a loading means, coupled to the missing frame window storing means, for loading the missing frame window storing means with a missing frame window associated with the sequence of the last received frame.
- 3. The interface device of claim 1, wherein the control circuit transfers data received by the interface device within a select frame, only after determining from the missing frame window that the select frame is the next expected frame.

- 25 -

4. The interface device of claim 1, wherein the control circuit transfers data received by the interface device within a select frame, only after determining from the comparison means and the missing window frame that the select frame is currently the next expected frame and the select frame was associated with a value within the predetermined range of values at the time the select frame was received.
5. The interface device of claim 1, wherein: (1) the control circuit has value signal lines that output the binary value associated with the last received frame, (2) the counter has output signal lines that output the binary value of the counter, and (3) the comparator window circuit further includes:
 - (a) a first comparator having an output and a first and second input, the first input coupled to the output signal lines of the counter, the second input coupled to the value signal lines, and the output of the comparator coupled to the control circuit, the first comparator for determining when the value associated with the last received frame is within a predetermined range of values relative to the value of the counter;
 - (b) an adder having an input and an output, the input coupled to the output signal lines of the counter, for outputting a value that is greater than the value of the counter; and
 - (c) a second comparator, having an output and a first and a second input, the first input coupled to the output of the adder, the second input coupled to the output signal lines of the counter, and the output of the second comparator being coupled to the control circuit, the second comparator for comparing the output of the adder to the value of the counter and providing the result of the comparison to the control circuit;

- 26 -

- (d) a subtractor having an output, a first input, and a second input, the first input coupled to at least one of the least significant output lines from the counter, and the second input coupled to at least one of the least significant output lines from the control circuit, for representing in binary form, the difference between the value represented on the least significant lines of the counter and the value represented on the least significant lines from the control circuit;
 - (e) a decoder having an input and an output, the output including a plurality of output signal lines, each output signal line being associated with a unique binary input value coupled to the decoder from the subtractor; and
 - (f) a plurality of two-input OR gates, each OR gate being associated with a unique one of the decoder output signal lines, a first input of each OR gate being coupled to the associated decoder output signal line, the output of each OR gate being coupled to a unique input to the missing frame window storing means for storing values indicative of the values associated with frames received by the interface device.
6. The interface device of claim 1, wherein the comparator window circuit further includes:
- (a) means for determining whether a predetermined number of high order bits of the value associated with the last received frame are equal to or one greater than an equal number of high order bits of the value associated with the next expected frame;
 - (b) means for asserting one of a plurality of difference signal lines, each difference signal line associated with a unique value, when the difference between the least significant bits of the value associated with the last received frame and an equal number of the least significant bits of the value associated with

- 27 -

the next expected frame is equal to the unique value associated with the difference signal line to be asserted; and

- (c) a plurality of two-input OR gates, each OR gate being associated with a unique one of the plurality of difference signal lines, a first input of each OR gate being coupled to the associated difference signal line, the output of each OR gate being coupled to a unique input to the missing frame window storing means for storing values indicative of the values associated with frames received by the interface device, each value being within a predetermined range of values relative to the next expected frame, the second input of each OR gate being coupled to an output from the missing frame window storing means.

1 / 3





FIGURE 1

2 / 3



3 / 3




FIGURE 3

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US96/03970

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : H04Q 11/04

US CL : 370/61

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/13, 18, 60, 60.1, 61, 82, 94.1

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
TANENBAUM, COMPUTER NETWORKS (2ND EDITION), TK5105.5 T36, 1988

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS, STN

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US, A, 4,654,890 (HASEGAWA ET AL) 31 March 1987	1-10
X	TANENBAUM, Andrew S., Computer Networks (2nd Edition), TK5105.5 T36, Prentice Hall, 1988, p. 136-184.	1-10

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be part of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"G" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

14 MAY 1996

Date of mailing of the international search report

24 JUN 1996

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

MELISSA KAY CARMAN

Telephone No. (703) 308-7605

VI. Example Implementations

The present invention (i.e., CIDM system 100 or any part thereof) may be implemented using hardware, software or a combination thereof and may be implemented in one or more computer systems or other processing systems. In fact, in one embodiment, the invention is directed toward one or more computer systems capable of carrying out the functionality described herein. An example of a computer system 2900 is shown in FIG. 29. The computer system 2900 includes one or more processors, such as processor 2904. The processor 2904 is connected to a communication infrastructure 2906 (e.g., a communications bus, cross-over bar, or network). Various software embodiments are described in terms of this exemplary computer system. After reading this description, it will become apparent to a person skilled in the relevant art(s) how to implement the invention using other computer systems and/or computer architectures.

Computer system 2900 can include a display interface 2905 that forwards graphics, text, and other data from the communication infrastructure 2906 (or from a frame buffer not shown) for display on the display unit 2930.

Computer system 2900 also includes a main memory 2908, preferably random access memory (RAM), and may also include a secondary memory 2910. The secondary memory 2910 may include, for example, a hard disk drive 2912 and/or a removable storage drive 2914, representing a floppy disk drive, a magnetic tape drive, an optical disk drive, etc. The removable storage drive 2914 reads from and/or writes to a removable storage unit 2918 in a well known manner. Removable storage unit 2918, represents a floppy disk, magnetic tape, optical disk, etc. which is read by and written to by removable storage drive 2914. As will be appreciated, the removable storage unit 2918 includes a computer usable storage medium having stored therein computer software and/or data.

In alternative embodiments, secondary memory 2910 may include other similar means for allowing computer programs or other instructions to be loaded into computer system 2900. Such means may include, for example, a removable storage unit 2922 and an interface 2920. Examples of such may include a program cartridge and cartridge interface (such as that found in video game devices), a removable memory chip (such as an EPROM, or PROM) and associated socket, and other removable storage units 2922 and interfaces 2920 which allow software and data to be transferred from the removable storage unit 2922 to computer system 2900.

Computer system 2900 may also include a communications interface 2924. Communications interface 2924 allows software and data to be transferred between computer system 2900 and external devices. Examples of communications interface 2924 may include a modem, a network interface (such as an Ethernet card), a communications port, a PCMCIA slot and card, etc. Software and data transferred via communications

interface 2924 are in the form of signals 2928 which may be electronic, electromagnetic, optical or other signals capable of being received by communications interface 2924. These signals 2928 are provided to communications interface 2924 via a communications path (i.e., channel) 2926. This channel 2926 carries signals 2928 and may be implemented using wire or cable, fiber optics, a phone line, a cellular phone link, an RF link and other communications channels.

In this document, the terms "computer program medium" and "computer usable medium" are used to generally refer to media such as removable storage drive 2914, a hard disk installed in hard disk drive 2912, and signals 2928. These computer program products are means for providing software to computer system 2900. The invention is directed to such computer program products.

Computer programs (also called computer control logic) are stored in main memory 2908 and/or secondary memory 2910. Computer programs may also be received via communications interface 2924. Such computer programs, when executed, enable the computer system 2900 to perform the features of the present invention as discussed herein. In particular, the computer programs, when executed, enable the processor 2904 to perform the features of the present invention. Accordingly, such computer programs represent controllers of the computer system 2900.

In an embodiment where the invention is implemented using software, the software may be stored in a computer program product and loaded into computer system 2900 using removable storage drive 2914, hard drive 2912 or communications interface 2924. The control logic (software), when executed by the processor 2904, causes the processor 2904 to perform the functions of the invention as described herein.

In another embodiment, the invention is implemented primarily in hardware using, for example, hardware components such as application specific integrated circuits (ASICs). Implementation of the hardware state machine so as to perform the functions described herein will be apparent to persons skilled in the relevant art(s).

In yet another embodiment, the invention is implemented using a combination of both hardware and software.

THIS PAGE BLANK (U8PTD)